



DOE's Advanced Network Initiative (ANI) Testbed Project Update

Brian Tierney, Staff Scientist

ESnet

January 12, 2012



Overview

Testbed Value

Testbed Overview

Current Testbed Research and Results

Testbed Future





Testbed Value

- A realistic **national-scale** environment for innovative network research that would be impossible for most researchers to create in their labs
- Many types of network research require realistic high-latency environments
- Maximum flexibility: researchers get “super-user” access to everything
 - ability to install custom OS on hosts
 - ability to do full routers/switch configuration
 - ability to install custom router software
- A controlled environment that supports reproducible results

Priority is given to research that aligns with ESnet strategic focus areas, such as:

- High performance network protocols, data transfer middleware, Openflow + MPLS integration strategies, energy efficient networking

Testbed Components



The ANI Testbed is really 3 separate Testbeds, each focused on different types of research

- Control Plane Testbed
- 100G Testbed
- Dark Fiber Testbed

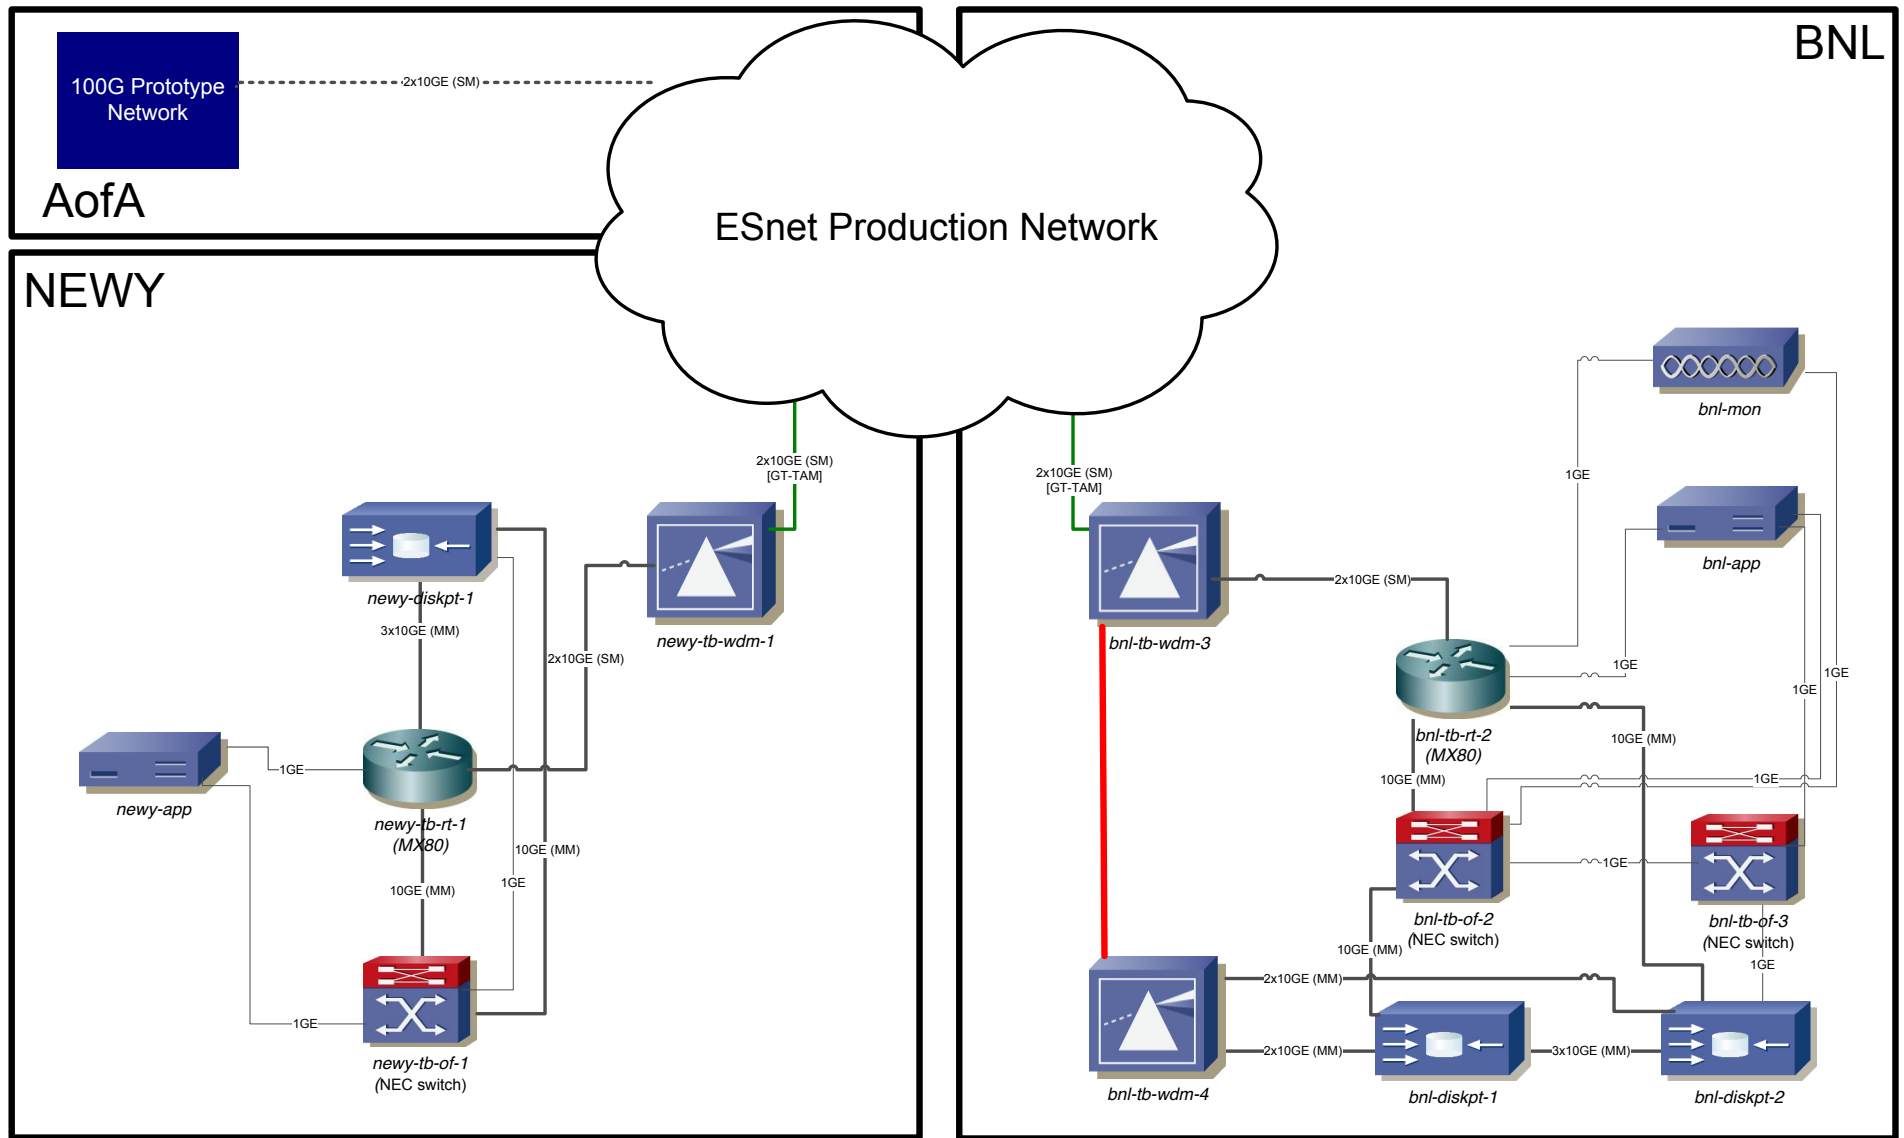


Control Plane Testbed

Control Plane Testbed

- 700 KM loop on Long Island Between Manhattan and Brookhaven National Lab
- Infinera DWDM switches, NEC OpenFlow switches, Juniper router (soon with OpenFlow support), 6 hosts
- Research Examples
 - OpenFlow experiments
 - Path Computation algorithms
 - MPLS/OpenFlow Integration

Long Island MAN (LIMAN) Testbed Architecture



Updated July 8, 2011

Notes:

- “App Host”: can be used for researcher application, control plane control software, etc. Can support up to 8 simultaneous VMs
- “I/O Testers” are capable of 15 G disk-to-disk or 35G memory-to-memory
- Other infrastructure not shown: VPN Server, file server (NFS, webdav, svn, etc.)



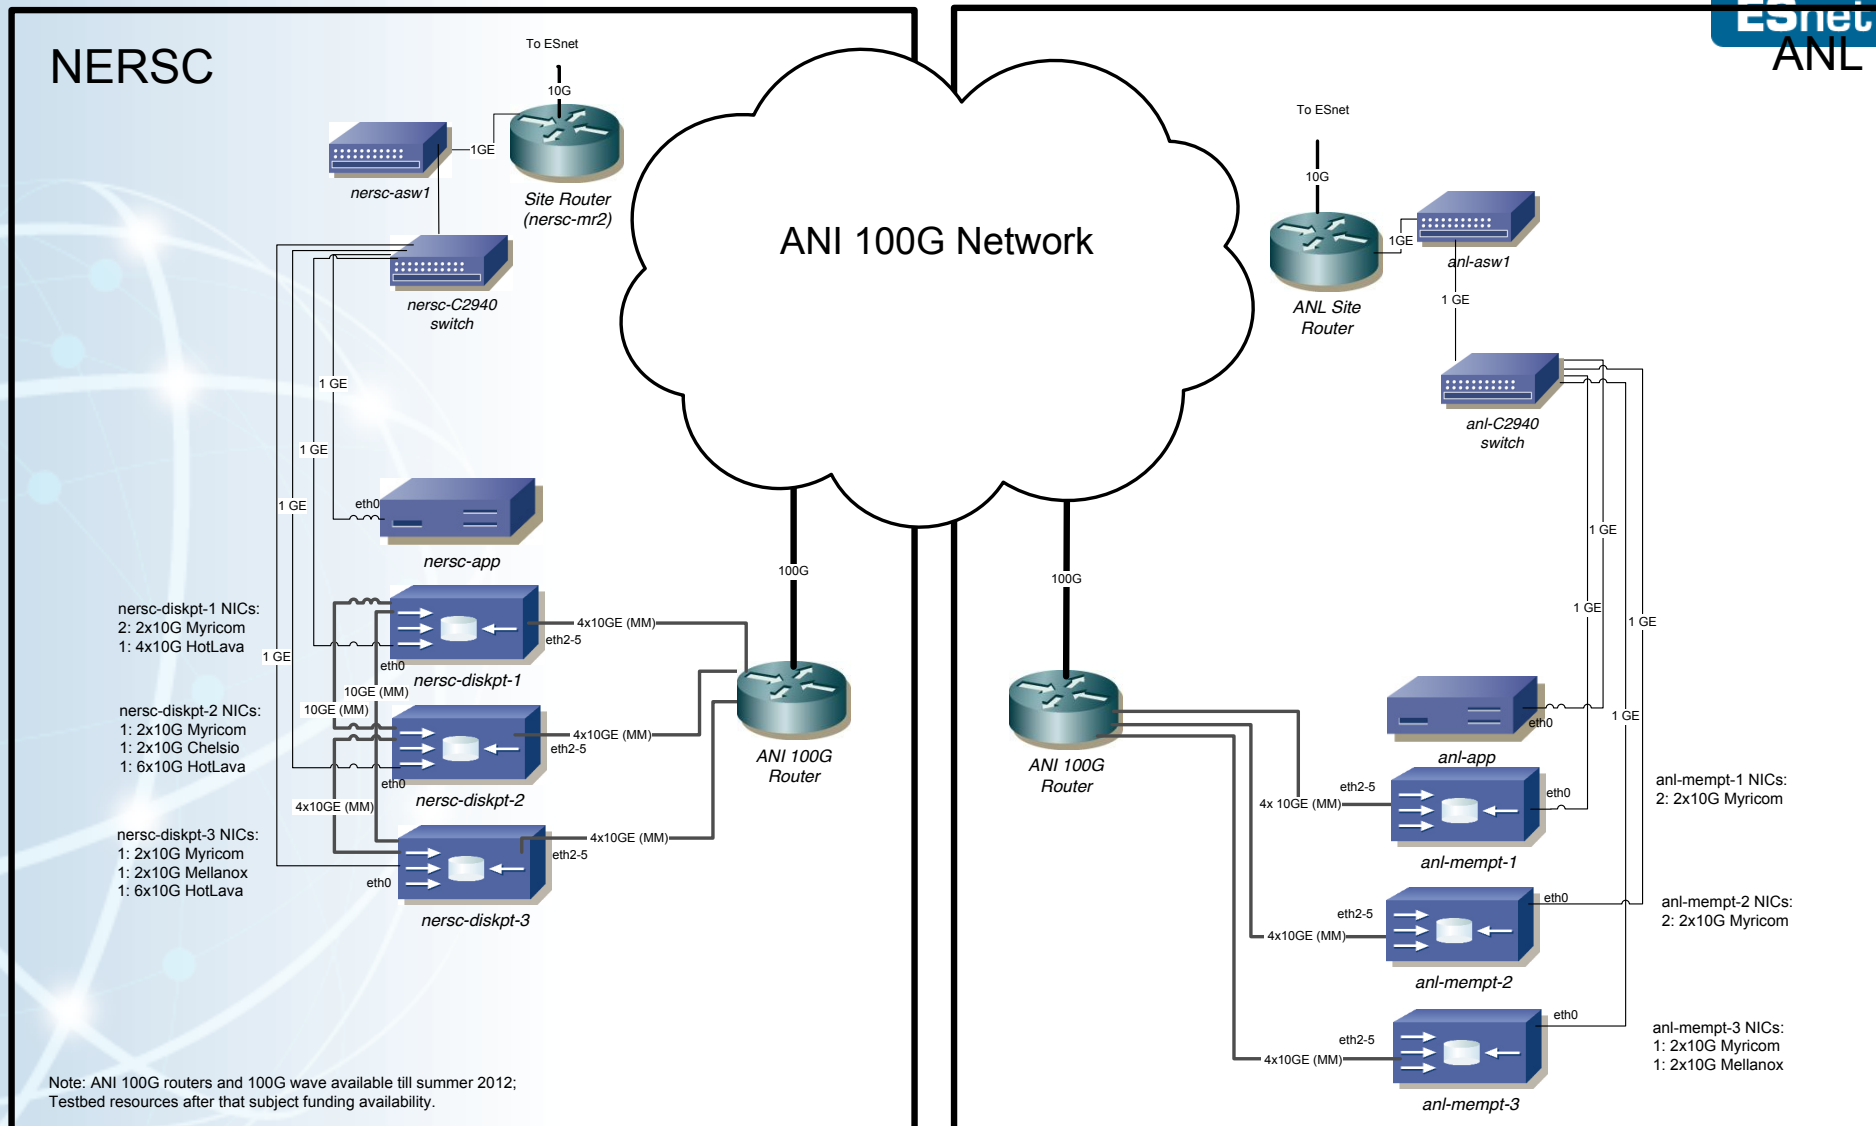
100G Testbed

100G Testbed

- 100Gbps path from Oakland, CA (NERSC) to Argonne, IL (ANL)
- 3 high-end hosts at each site, capable of >39 Gbps each
- Research Examples
 - 100G data transfer tools
 - FTP100, GridFTP, BestMan, Saratoga
 - TCP alternatives
 - RDMA, RoCE
 - TCP enhancements
 - PERT, RAPID TCP

Available as of Jan 3, 2012

ANI Middleware Testbed



1/12/12

Updated December 11, 2011
8



Related Resources

Magellan at NERSC

8 nodes with 2x10GE, GPFS backend

Magellan at ANL

16 nodes, but going away very soon

ORNL test cluster

12 nodes, LUSTRE backend

FNAL test cluster

in planning stage



Dark Fiber Testbed

Dark Fiber Testbed

- Nationwide footprint (13,000 miles)
- Researchers must pay to light the fiber, but no charge for access if results will be public
- Research Examples
 - Greentouch Energy experiments
 - JPL quantum networking / quantum cryptography experiments

Dark Fiber Network

LBNL Long Haul Dark Fiber Routes
12,924 miles



1/12/12

11



Current ANI Testbed Research

25 projects have been accepted to the testbed so far:

- 8 via direct DOE/ASCR funding

- 17 via testbed proposal review process

- 5 from Industry; 8 from DOE labs; 3 from NASA, 1 from DISA, 7 from Universities

Wide range of projects:

- 9: high-speed middleware

- 3: OpenFlow

- 4: other network control plane

- 1: 100Gbps end host hardware

- 1: network flow classification

- 1: quantum communications

- 2: Wide Area RDMA

- 2: TCP congestion control

- 1: security

- 1: energy efficiency

<http://www.es.net/RandD/advanced-networking-initiative/current-testbed-research/>

1/12/12

12

Potential Breakthroughs from Current Testbed Research



Unconditionally secure continuous variable quantum communications

Ability to scale TCP well beyond 10Gbps

Alternative transport protocols that scale better than TCP (e.g.: RDMA)

Ability to easily create end-to-end circuits

Ability to scale data transfer tools/middleware to 100Gbps and beyond

100Gbps host interface

For more information see:

<http://www.es.net/RandD/advanced-networking-initiative/current-testbed-research/>

Results to Date: Publications Submitted



After only 7 months of being operational, there are already several publications submitted based on ESnet's Testbed results:

- *Energy Profiling of Network Elements for Rate Adaptation* Technologies, M. Ricca, A. Francini, S. Fortune, T. Klein, submitted to the ACM/IEEE 3rd International Conference on Future Energy Systems (e-Energy 2012).
- *Traffic engineering in hybrid IP/optical circuit networks*. Zhenzhen Yan, Chris Tracy, and Malathi Veeraraghavan, submitted to the IEEE 13th Conference on High Performance Switching and Routing (HPSR)
- *Middleware Support for RDMA-based Data Transfer in Cloud Computing*, Yufei Ren, Tan Li, Dantong Yu, Shudong Jin, and Thomas Robertazzi, submitted to IEEE International Parallel & Distributed Processing Symposium (IPDPS)
- *Identifying gaps in Grid middleware on fast networks with the Advanced Network Initiative*, G. Garzolglio and Dykstra, accepted at International Conference on Computing in High Energy and Nuclear Physics (CHEP 2012)

Results to Date: Press Releases



Joint Press Release by Orange Silicon Valley, Bay Microsystems and ESnet on the first wide area 40G RDMA results (<http://goo.gl/57gSJ>)

- Showed that RDMA moves data at up to 96 percent of the peak capacity of the network
 - Much more efficient than TCP
- Replicated at SC11 over a 10,000KM path (<http://goo.gl/ZE0w8>)

Greentouch press release on plans to use ESnet dark fiber testbed for energy experiments: <http://goo.gl/OB6nD>

Results to Date: Demonstrations



- ESnet Demonstration of OSCARS/OpenFlow integration on the ANI testbed at the Open Networking Summit held at Stanford.
 - <http://goo.gl/4FFBg>
- ESnet Demonstration of 9.9Gbps RDMA transfers from the ANI Testbed at BNL to Seattle at SC11

Results to Date: Industry Involvement



- Loaner hardware for 40GE testing
 - Infinera, Alcatel-Lucent (ALU), Bay Microsystems, Mellanox



Future Testbed Challenges

Long term operational support of the Testbed

- Current testbed funding will run out in August of this year
- Currently funded at 2FTE for testbed support
- This is essential for continued success

Hope to secure funding to ensure persistent/indefinite availability

- 100G wave from NERSC to ANL will transition to production use starting this summer, so need to identify funds to light a 2nd wave for testbed use

Dark Fiber will be available to the testbed until it is needed for production use

- For some segments this may not be for several years.

More Information



<http://sites.google.com/a/lbl.gov/ani-testbed/>

email: **ani-testbed-proposal@es.net**, BLTierney@es.net